

Digital Subscriber Line Network Deployment Method

Related Applications

The present application claims the benefit of U.S. Provisional Application Number 60/191,797 filed on March 24, 2000 and entitled "Digital Subscriber Line
5 Deployment Method".

Field Of The Invention

This invention generally relates to methods and systems that perform network planning and in particular such method and systems that are used to design networks and deploy equipment in such networks, in particular broadband access networks.

Background Of The Invention

10 Demand for Internet access and convergent broadband services has required traditional or incumbent service providers to evolve into high-speed Internet and multimedia providers. One of the main drivers of this evolution or change is Digital Subscriber Line (DSL) technology. DSL has allowed service providers, both
15 incumbent and competitive, to create approaches that address this demand, or opportunity, and has been serving as a key enabler of this paradigm shift. As a result, deployment of DSL products has been doubling every six months. As this momentum continues to build, with growth projections of more than 85% per year, efficiency and cost containment are becoming critical. In order to capitalize on this marketplace,
20 service providers must be able to quickly forecast customer demand and deploy their resources as effectively as possible. In short, the need for superior planning is acute.

As the demand for broadband capability increases, network providers are faced with decisions on which access architectures to deploy in which service areas. The choice of DSL technologies varies based on network geography, customer service

demand, pre-existing structures, and the specific design rules associated with each technology option.

DSL technologies work by connecting a pair of “DSL modems” to a subscriber’s existing telephone line. One of these modems resides at the subscriber’s premises, while the other is built into a line card at a Digital Subscriber Line Access Multiplexer (DSLAM) placed at either the central office or an appropriate cross-connect location in the field. Since subscriber premise modems are dedicated to each subscriber, their cost represents a fixed cost for DSL deployment. In contrast, DSLAMs are shared among subscribers and are expensive, so placing them judiciously can result in cost savings. DSLAMs are available in a number of different configurations that serve different numbers of subscribers at different cost levels. The variety of options and the potentially large number of subscribers make cost-effective placement of DSLAMs a complex problem.

In order to better understand the underlying problem, DSL may be viewed as being overlaid on an existing copper network within a single wire center. The existing copper network defines how subscribers are connected to the central office and dictates where DSLAMs can be placed. This network is logically a tree connecting subscribers to the central office. Typically, DSLAMs can be placed at the central office and suitable cross-connect locations in the field. However, in the context of long-range network planning, future subscribers to DSL are unknown, so forecasts can be a key driver for early deployment decisions.

The choice of DSLAM locations for a particular subscriber is further restricted by range constraints that impose limits on the allowable wire length from the subscriber to the serving DSLAM. This range limit is typically dictated by the type of

DSL service, e.g., Asymmetric DSL (ADSL), Hybrid DSL (HDSL), Very high bit-rate DSL (VDSL), etc., to be deployed and the physical properties of the copper plant.

In general, prior art approaches to determining the location to deploy or place DSLAMs are manual in nature, requiring an operator to create a network architecture that will best meet future customer demand based on manual design rules. In particular, the demand models used by an operator suffer from the major shortcoming that the model is not directly related to potential short-term or long-term customer demand within a particular area. For example, a prior art approach comprises placing DSLAMs at the hubs of a hub and spoke model. By this deployment model the hubs are the host locations and correspond to either a corporate Local Area Network (LAN) or Internet Service Provider location. The subscribers are then connected via the spokes to the host location. This method or model for deployment is under-inclusive in that there may be many subscribers “beyond the reach of the host/hub” who desire DSL service. In addition, this method unnecessarily ties demand to the anticipated needs of large corporations and ISPs (Internet Service Providers). More importantly, deploying equipment in accordance with this methodology does not take into account future growth both within and outside the reach of the hub and is in no way related to a customer demand forecast.

Of utility then are methods and systems for placing and allocating DSLAM capacity to serve current and future demand for DSL service in a geographic area and which methods and systems are based on customer demand forecasts.

Summary

Our invention is a method and system for placing DSLAM equipment to serve estimated demand for DSL service within a geographic area. Although our method is described below in the context of DSLAM deployment, it is also generally applicable

to the deployment of other network equipment or architectures, e.g., Fiber To The Curb (FTTC) or other FTTx equipment or architectures.

In general our method begins with the selection of a planning area. The planning area may include a central office, cross-connects, copper facilities and subscriber locations within the planning area. The planning area may be conveniently
5 thought of as comprising a baseline network or architecture having a tree structure wherein the copper facilities comprise the paths or branches interconnecting the cross-connects, central office, and subscriber locations; in addition, the cross-connects and central office comprise the nodes of the tree and serve as possible sites
10 where DSLAMS may be placed. Based on the chosen planning area or existing baseline architecture a service demand model or demand forecast model based on socio-demographic data is used to generate a demand forecast for the area comprising the baseline network or its constituent sub-areas. The demographically-driven forecast is then used to generate a deployment plan or a demographically-driven
15 network architecture wherein the DSLAMs are placed at selected nodes or sites in accordance with certain design constraints while minimizing deployment costs.

In accordance with an aspect of our invention, the demographics of the planning area are taken into account in determining DSL service demand within the planning area. That is, demographic data accessible from a geographic information
20 system (GIS) database is used as input to the demand forecast model. In accordance with this aspect of our invention the planning area is linked to a geographic information system to produce a demographically-driven demand forecast model. As such our invention is able to evaluate DSL options and select the best option to deploy in a particular area.

The demographically-driven demand model provides input in determining the sites wherein DSLAM equipment may be placed while minimizing cost and satisfying other design constraints. The sites for DSLAM placement are then displayed as a network tree wherein each site is interconnected. The network tree that is generated includes identification of the sites at which the equipment is to be located and for each site or node the number and type of equipment that needs to be deployed.

In accordance with an object of our invention, demographic data is advantageously accessed from a global information system (GIS) database and used to create a demographically-driven demand model. By utilizing the GIS database our invention is better able to estimate the demand within an area. Further in accordance with our invention, our method for developing a demographically-driven demand model includes determination of the number of households in each neighborhood block estimated to buy service.

Our method therefore advantageously allows a service provider to deploy equipment in a manner that meets near-term demand while at the same time providing for future growth within an area. In addition, our method is not under-inclusive and depends more on consumer's ability to pay and other characteristics that make it more or less probable that consumers within an area will be DSL service subscribers.

Description of the Drawings

FIG. 1 is a flow chart of the method steps for deploying network equipment in accordance with our invention;

FIG. 2A depicts the steps for selecting a planning area in accordance with the present invention;

FIG. 2B depicts an exemplary planning area and accompanying baseline network architecture of the copper cable facilities within the exemplary planning area;

FIG. 3 depicts the steps for determining customer demand in accordance with the present invention;

FIG. 4 depicts the steps for generating a demographically-driven deployment or network architecture in accordance with the present invention; and

5 FIG. 5 is an exemplary network tree diagram having candidate DSLAM sites for equipment placement in accordance with the present invention.

Detailed Description

Turning to FIG. 1 there is illustrated a flow chart of the method steps of our invention. As FIG. 1 shows, the method begins with selecting or defining a planning
10 area, block 1100. The selected planning area is then used to determine customer demand based on a demographically-driven service demand model and forecast, block 1200. The demographically-driven demand forecast is then used as input to generate a deployment plan or layout model/network architecture, block 1400. The demographically-driven demand plan or network architecture places equipment in the
15 network, block 1500, at candidate sites.

Turning now to FIG. 2A there is illustrated the substeps, blocks 1105 to 1129, for defining a planning area (block 1100 of FIG. 1). As FIG. 2A shows, the first substep of defining a planning area is selecting a geographic area, block 1105. In general, and in accordance with our invention, any type of spatial unit that a network
20 planner would find useful can be used in selecting the geographic area, block 1106. For example, a geographic area could be selected based on any customer or distribution area 1107, a wire center 1109, or a specific cable facility within a wire center 1115. The selected geographic area essentially defines the scope of the deployment problem. Choosing an entire wire center creates a larger and more
25 difficult problem because a wire center typically covers a fairly large geographic area

having many customers. Accordingly, less computing resources are required if the geographic area is limited to a cable within a wire center 1115. On the other hand, it should be noted that the deployment plan or network that is the output of our invention changes depending on the input. Therefore, although more difficult

5 computationally, it may be more efficient, with respect to cost, to define an area having an entire wire center comprising the planning area. Moreover, the accuracy of the demand forecasts is improved by increasing the number of sampling units, hence demographic variability, in the planning area.

In accordance with an embodiment of our invention, once the geographic area

10 is selected, block 1105, a geographic information system (GIS) may be accessed, block 1125, to ascertain demographic information on the subscribers within the selected geographic area. As those of ordinary skill in the art will know, the GIS is a very useful visual aid that can be used, and was used in our implementation, to select the geographic area, 1105. The information from the GIS database may be used to

15 determine the boundaries of the selected geographic area, block 1127, and to link the geographic area to a wire center or centers, block 1129.

An exemplary display of wire centers linked to the geographic area based on the information accessed from the GIS database is shown in FIG. 2B. The database should include information on individual census block group boundaries 1150. As

20 discussed below, the information in the database may also include wire center map boundary information, the central office within a wire center, etc.

As FIG. 2B shows, the wire center information that is obtained from the GIS may be used to create a network tree 1160, having central office 1162 and cross-connects 1164, which are interconnected by the copper cables 1170 that also connect

25 the potential subscribers 1174 (more than one subscriber was not shown to avoid

clutter). This network tree 1160 represents a baseline network having many sites wherein DSLAM equipment may be placed, i.e., central office 1162 and cross-connect 1164, so as to provide DSL service to the potential subscribers who are also connected to the network tree. The problem therefore becomes one of placing the
 5 DSLAMs at the available sites, given customer demand and network design constraints, while minimizing deployment costs.

With the geographic area selected the next step in our method, returning to FIG. 1, is that of determining subscriber demand, block 1200. In accordance with our invention subscriber demand is demographically driven, as indicated by the steps
 10 illustrated in FIG. 3. That is, subscriber demand is determined from or based on demographic information available in a GIS.

With reference to FIG. 3, once the demographic data is available we then associate the demographic data with the planning area, block 3240. It should be noted that GIS databases are commercially available and we have found that such a database
 15 should include the following geo-referenced data layers or information:

- 1) Wire Center boundary information
- 2) Central Office locations within these Wire Centers
- 3) Map of copper network within these Wire Centers
- 4) Precise location of nodes where DSLAMS might be deployed
- 20 5) Cable connections joining nodes
- 6) Street map for wire centers (In vector format with feature files needed to identify Census Bureau within block groups)
- 7) US census block group boundaries
- 8) Thematic map layers Census Bureau's Block Groups contained within
 25 the study Wire Centers including

- a) Average household income
 - b) Average days/week work at home
 - c) Average number of telephone lines (if available)
 - d) Average commute distance to work or average number of
5 days/weeks work at home (if available)
 - e) Average educational attainment (if available)
 - f) Employment profile by job class (if available)
 - g) Proportion of single and multi-family residences
- 9) DSL service profile for the equipment being deployed .
- 10) Map of the Fiber network in these locations

The information or data listed above is meant to illustrate data that is useful in performing an accurate forecast. Accordingly, all of the above information need not be used to get an accurate forecast. On the other hand, depending on the population make-up of the area certain data will better predict demand than other data, as
15 indicated below. In addition, the above list of information is not meant to be exhaustive. In fact, there may be other information that may be useful in determining service demand.

In associating the GIS data, block 3240, relevant demographic data for each block group in the study area is selected, including a response variable such as
20 average household interest in high-speed service and a collection of predictor variables, e.g., average household income, average educational attainment of household head, average number of phone lines, etc. A user can select any of these variables to estimate the proportion of households in a block group that are likely to subscribe to DSL services or other broadband services.

An estimate of the households likely to subscribe, block 3250, is obtained, in accordance with our invention, by fitting a Mixed Regressive-Spatial-Autoregressive Logistic (MRSAL) model to the selected data, block 3251. Our MSRAL model exploits spatial patterning and temporal relationships to improve the forecast by exploiting spatial relationships. For example, if a census block group will subscribe to service then it is more likely that a neighboring census block group will also subscribe to service. Our model is of the following form:

$$P(W; x_1, x_2, \dots, x_k) = 1 / [1 + \exp\{-(\alpha + \rho WP + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_k x_k)\}],$$

where:

- 10 P = probability of a household subscribing to xDSL services;
- W = a spatial weights matrix whose values take on a nonzero value when census block groups are neighbors and a zero value otherwise;
- ρ = the spatial autoregressive coefficient or effect of spatial dependency on P while controlling for the effects of the x_i ;
- 15 x_i = predictor variable i ;
- k = number of predictor variables;
- α = constant (to be determined); and
- β_i = effect of predictor i on P (to be determined).

At the block (i.e., neighborhood block) level, all subscribers are assumed to home onto the network through the same cross-connection location, thereby eliminating the need for further desegregation. Demand is then the estimated number of DSL subscribers on a given block. The association between demographics in the planning area and the resulting estimate produce a demographically-driven service demand forecast. In accordance with an embodiment of our invention, the DSL forecasts for all block groups within a planning area may be displayed on the GIS display. Note,

however, that in accordance with our invention the important aspect is not the actual display of the demographically-driven information nor the form of the display, instead it is use of demographics to drive the demand model or forecast. Those of ordinary skill in the art will note that the idea here is to determine which variables are useful
 5 predictors, given those available, and use the GIS framework to estimate the demand forecasts needed by the placement algorithm.

In addition, in some cases use of neighborhood information may not be warranted or necessary when forecasting the demand. Neighborhood information is an additional term that is used to improve the precision of the forecast thus all the
 10 values of W may be set to zero making the term ρWP in the above expression equal to zero. In this case, the forecast is made only with the demographic data (the x_i 's) with no spatial effects.

In the end, a demand forecast is produced, block 3254, as result of application of our demand model. This demand forecast, block 3254, can be used to seed a GIS
 15 database, appearing in the database as datalayer.

With reference to FIG. 2B, the estimate may affect the network tree 1160 by indicating that certain branches 1176 on the tree may be removed because the forecasted demand for service in that census group is below a predetermined threshold. Specifically, it may not be profitable to service an area where only 30% of
 20 the census block group will desire service.

At block 1400, FIG. 1, the demand forecast obtained from the demand forecasting model above is used as an input to produce a deployment plan or demographically-driven network tree (or map) showing the actual location of the equipment and the capacity of each of the DSLAMs to be deployed; note, the GIS
 25 display may be used to visualize a map of the deployment plan having the actual

location of the newly placed DSLAM equipment. Turning to FIG. 4, there is depicted the steps for generating a deployment plan in accordance with the present invention. In order that the demographically-driven demand or deployment map may be generated, the length of the cable connecting each potential customer is normalized,

5 block 4310. Cable normalization 4310 is necessary because the cable, or subscriber loop, connecting any given customer may be made of several different gauges, e.g., 26 or 24 gauge cable, with different attenuation properties. That is, cable gauges may differ from one customer to another. At block 4320, using the demand for each block a determination is made of how each block homes onto the existing baseline network.

10 For example, it may be assumed that each subscriber in a census block homes onto the baseline network tree through the same cross-connect location or site. At block 4330, the available locations for placing the equipment, along with the cost, capacity, and maximum reach of the equipment is determined. At block 4340, a tree network connecting all the available locations or sites for the equipment is constructed. Using

15 the results from blocks 4310 – 4340 as input, optimization is performed, block 4350, to place the DSLAMs at the available or designated sites and customers are assigned to serving locations so that the total cost of the DSLAMs is minimized.

The output from the optimization process provides two types of information. For each demand location, the site at which its serving DSLAM is located is

20 identified; these sites may be conveniently referred to as network nodes. For each network node, the set of DSLAMs (possibly none) to place is specified. With this information, the total deployment cost is computed and includes: cost of modems; cost of DSLAMs; and since the DSLAM locations define where the cutover between fiber and copper occurs, the cost of fiber may also be included. The configuration of

25 the DSLAM is arbitrary and can be defined by the user, i.e., DSLAMs have a

common housing and capacity is installed in modular components. In implementing the present invention we chose four "virtual" DSLAMs, at 25%, 50%, 75%, and 100% utilization. Note however, that our method can accommodate any set of virtual DSLAMs.

5 For example, if a network or service provider desires to determine where DSLAMs should be placed in a geographic area to support Asymmetric Digital Subscriber Line (ADSL) technology, the optimization method or optimizer, block 4350, would operate as follows. The optimizer takes as input: the set of DSLAMs available to place; the "tree" network (refer to FIG. 2B) that connects all of the
10 candidate DSLAM sites (one of which is the central office 1162 of FIG. 2B); customers with their associated homing sites; and the maximum allowable distance between a customer and its serving DSLAM. Given this information, the optimization method places DSLAMs at designated sites, 1164 FIG. 2B, and assigns customers to serving locations so that the total cost of the DSLAMs is minimized and
15 so that certain design constraints are satisfied. For the ADSL design problem these constraints assure that: 1) no customer is too far from its serving DSLAM; 2) no DSLAM serves more customers than its allowed capacity; 3) all customers are served at a site along the path to the central office; and 4) any two customers whose copper wires meet on their path to their serving DSLAM must be served at the same DSLAM
20 location.

The constraint on the length of copper to the DSLAM ensures that a customer receives sufficient signal strength and quality. The second constraint is a capacity constraint at the DSLAM. The third constraint arises because ADSL is typically deployed on top of an existing copper loop. By assigning customers to be served at
25 DSLAM locations along the path to the central office or root, we guarantee that no

additional copper needs to be installed. Note however, that the constraint that a subscriber must be served along the path to the central office is a constraint owing to overlaying the new network onto the existing plant. This constraint could be relaxed where newer network types, such as FTTC, are being considered for deployment.

5 Finally, the fourth constraint is something of a regularity constraint that seems to produce designs that "look right" to network designers and will be easier to maintain. This constraint is a pivotal assumption underlying the optimization algorithm.

The optimization method is a "dynamic programming" (DP) method. It begins at the leaves of the tree formed by the DSLAM locations and works from the leaves to

10 the central office.

In general, dynamic programming is a method for solving optimization problems that have a large search space but are structured in such a way that it is possible to "grow" a full solution from optimal partial solutions. A common thread in dynamic programming is solving a complex problem by solving a sequence of simpler

15 problems. In ADSL design or any DSL design (HDSL, SDSL, VDSL), we begin by solving small problems at the leaves of the tree. At a particular leaf, we must decide whether to serve the customers that home there with a DSLAM at that node or at some node closer to the central office. If we decide to serve the customers at that leaf, then we must place a DSLAM there, so we'd choose the cheapest DSLAM (or

20 combination of DSLAMs) that can serve the subtending customers. If we decide to serve these customers closer to the central office, then the cost at the current leaf is zero, but these customers must percolate up the tree (toward the central office) until they are served by some DSLAM.

Typically, solving a dynamic program involves a certain amount of book-

25 keeping that is handled by defining an "optimal-value function". The optimal-value

function maps (in this case) decisions into costs. A key property of a the optimal-value function is that we can grow solutions by manipulating optimal-value functions.

To describe our method, it suffices to describe its optimal-value function. We will do that after we define a few helpful terms. First, and by way of reference to FIG.

5 5, a "node" 510 refers to any potential DSLAM location (including a central office 515). The "parent" of any node is the first node on the path between it and the central office. The "children" of a node are the ones that have that node as a parent. For any given node, we refer to its "induced subtree" as being all nodes and customers (including the node itself) that must cross this node to reach the central office.
10 Finally, the direction "up the tree" is from a node toward the central office and "down the tree" is from a node toward the leaves 516.

At a high level, an optimal-value function is computed at every node and it tells us how much it will cost to serve the customers in that node's subtree when they are served in a particular way. We compute the optimal-value function for a parent
15 node directly from the optimal-value functions of its children.

For the current problem, the optimal-value function is rather complicated. The optimal-value function is $c(P,V,n)$ which represents the least cost of serving all of the customers in node P's subtree except for n of them that are served at node V, where V is a node on P's path to the central office.

20 The method begins at the leaves of the tree (i.e., the nodes with no children) and proceeds up the tree, computing the optimal-value function for non-leaves only after the functions for their children have been computed.

To illustrate and with reference to FIG. 5, we'll provide a very small example. Suppose, that we have DSLAMs with the following costs and capacities:

25

COST	CAPACITY
1000	200

1500 400

Now suppose that we have nodes A and B that are children of node C and that there are 300 customers homed to A, and 200 homed to B and C. We'll let X denote any particular ancestor of C that all customers homed to A, B, and C can reach. For node

- 5 A: We must serve either all or none of the houses that home at A at node A. Thus, we either serve all 300 at A, or we serve them at node C or somewhere above C.

$$\begin{aligned}
 &c(A,C,300) = 0 \quad (\text{The cost at A is 0 because we don't need a DSLAM}) \\
 &c(A,X,300) = 0 \quad (\text{The cost at A is 0 because we don't need a DSLAM}) \\
 10 \quad &c(A,C,0) = 1500 \quad (\text{We need a DSLAM at A}) \\
 &c(A,X,0) = 1500 \quad (\text{We need a DSLAM at A})
 \end{aligned}$$

The case at node B is analogous:

$$\begin{aligned}
 &c(B,C,200) = 0 \\
 15 \quad &c(B,X,200) = 0 \\
 &c(B,C,0) = 1000 \\
 &c(B,X,0) = 1000
 \end{aligned}$$

Now, the case at node C must combine information from its children. Note that if a DSLAM is placed at C, then customers homed to A and B must be served at or below

- 20 C. This somewhat simplifies the calculation below.

$$\begin{aligned}
 &c(C,X,700) = 0 \quad (\text{The cost at C is 0 because custs are served at X}) \\
 &c(C,X,500) = 1000 \quad (\text{which is } c(B,C,0) + c(A,X,300)) \\
 &c(C,X,400) = 1500 \quad (\text{which is } c(A,C,0) + c(B,X,200)) \\
 25 \quad &c(C,X,200) = 2500 \quad (\text{which is } c(A,C,0) + c(B,C,0)) \\
 &c(C,X,0) = 3000 \quad (\text{which is the cheapest combination of DSLAMs to} \\
 &\quad \text{serve all customers at C}) \\
 &\quad \text{Note that this is less than:} \\
 &\quad c(A,C,0) + c(B,C,0) + \text{cost to serve 200 at C}
 \end{aligned}$$

- 30 If C were in fact the central office, we'd see that the cheapest way to serve the given set of customers would be \$3000. Needless to say, a realistic-sized problem could require very large structures to store the optimal-value information.

The above description has been presented only to illustrate and describe the invention. It is not intended to be exhaustive or to limit the invention to any precise

form disclosed. Many modifications and variations are possible in light of the above teaching. The applications described were chosen and described in order to best explain the principles of the invention and its practical application to enable others skilled in the art to best utilize the invention on various applications and with various
5 modifications as are suited to the particular use contemplated. Specifically, although the above description illustratively uses digital subscriber line technology those of ordinary skill in the art will note that our invention may be extended to the placement of capacity at the nodes of a tree-shaped network.